

Real-Time Multipath Mitigation with Sensor-Aided Long Coherent Integration (SALI)

Zhenlan Cheng, Maxim Köhler, Alessandro Bionson, Robert Lluís Garcia, Gregor Dumphart,
Konstantinos Arkoudogiannis, Christian Bischof, Duarte Dias, Sebastian Carreno
U-blox AG, Thalwil, CH-8800, Switzerland

BIOGRAPHY

Zhenlan Cheng holds a Ph.D. in Electrical Engineering, specialized in communications; crafting high-precision, low-power GNSS receivers since 2008 for mass markets; leads U-blox's Measurement Engine teams.

Maxim Köhler holds an M.Sc. in Electrical Engineering and Information Technology; developing positioning algorithms for mass market GNSS receivers since 2017; leads as U-blox's Positioning Engine architect.

Alessandro Bionson holds a Ph.D. in Telecommunications Engineering, specialized in signal processing. His research focuses on optimization problems and the development of algorithms for positioning technologies.

ABSTRACT

We present a novel multipath mitigation scheme, based on Sensor-Aided Long coherent Integration (SALI), for Sensor Fusion (SF) Global Navigation Satellite System (GNSS) receivers (RXs). Applied to terrestrial mobile channels in the GNSS downlink, the upgraded SF RX significantly outperforms the conventional SF RX based on ≤ 20 ms Coherent Integration (CI). The benefits are twofold: higher line-of-sight measurement accuracy in the presence of multipaths and enhanced immunity against non-line-of-sight measurements. Both are achieved with a marked drop of overall computational complexity. An embedded implementation of the SALI-based positioning and measurement engine on a single chip with conventional GNSS accelerators and an ARM Cortex M3 processor is proven feasible. In extensive road tests, conducted in open sky, foliage, suburban and deep urban environments, we analyze the pseudo-range and Doppler errors as well as position fix performance in terms of accuracy, integrity and availability w.r.t. the required 2m accuracy using L1 signals. In addition, we compare those statistics for different carrier-to-noise density ratio (C/N0) and for different rover speeds. We demonstrate that the improved measurement quality translates into substantially enhanced position solutions in all challenging environments. The SALI-based SF RX can run real-time CI of ~ 1 s, perfectly aligned with 1Hz GNSS-based navigation rate. This, in turn, allows for enhanced real-time position output in high-rate such as 100Hz, using sensor-based propagation. In the absence of sensor-aided motion compensation, we introduce a backup solution that runs 100ms CI without estimating Doppler rates explicitly. The shorter snapshots drive a Kalman filter that exploits the inherent motion dynamics. Even here, road test results are considerably improved in multipath rich environments with significant frequency dispersion.

INTRODUCTION

Multipath (MP) interference poses a substantial challenge for GNSS RXs, impairing positioning accuracy and integrity. A significant amount of research effort has been dedicated to signal processing techniques that mitigate MP effects. Traditional non-parametric techniques that leverage the time-domain correlation include narrow correlators (Van Dierendonck et al. (1992)), strobe correlators (Garin et al. (1996)), Double-Delta (DD) discriminators (McGraw and Braasch (1999)), etc. These techniques are effective against MPs exhibiting substantial delay offsets w.r.t. the Line-Of-Sight (LOS). However, they encounter difficulties with near MPs possessing a delay offset close to the correlator tap spacing. Therefore, they are quite effective in the high C/N0 regime, where narrow tap spacing can be applied to fully exploit the capabilities of a large pre-correlation bandwidth without being compromised by code jitter. However, MP scenarios often coincide with a significant drop in C/N0, where time-variant MPs and/or interferers—in addition to noise—increasingly contribute to code jitter (Betz and Kolodziejewski (2009)). This challenges the robustness of the Delay-Locked Loop (DLL) and reduces the applicability of a high-resolution delay discriminator that relies on narrow tap spacing. The DLL performance is rather predictable in a quasi-stationary RX, where the DLL can converge to the set-point of the discriminator. In mobile applications, however, DLL's low-pass filtering effect becomes a critical factor when it interacts with time-variant MPs. A narrowband DLL, for instance, exhibits a degree of immunity against dynamic MPs whose relative frequencies to the LOS exceed the DLL bandwidth. This suggests a potential advantage in directly exploiting channel properties in the frequency domain, in addition to the delay domain, to suppress MPs before they reach the discriminator or estimator. On the other hand, we also recognize the benefits of the code tracking agility provided by a wideband DLL that facilitates swift re-convergence to the LOS when transitioning from a Non-Line-Of-Sight (NLOS) environment. Yet, adaptive DLL bandwidth with variable time-constants produces unpredictable correlation among consecutive measurements.

This can lead to underestimated position errors, posing challenges for integrity monitoring methods, especially in safety-critical applications where it becomes particularly problematic. These considerations add layers of complexity to the DLL design. In practice, achieving optimal performance in the presence of MPs requires skillfully balancing multiple design objectives to navigate the trade-offs inherent in real-world scenarios.

It's well-known that a Sensor Fusion (SF) architecture can substantially improve the position fix quality, particularly in GNSS denied environments, by incorporating an additional Inertial Measurement Unit (IMU) into the Positioning Engine (PE) (Schmidt and Phillips (2010); Schmidt and Phillips (2011); Somieski et al. (2010)). Upon proper sensor calibration, the velocity estimation can be enhanced by the IMU, leading to improved MP immunity. Moreover, the accurate velocity estimate can be mapped to LOS Doppler in real-time, making the local oscillator follow Doppler dynamics and largely wipe them off. This in turn improves the tracking performance of conventional Frequency-Locked Loops (FLL) and Phase-Locked Loops (PLL) because of reduced dynamics (Gao and Lachapelle (2008)). On the other hand, the recurring sensor calibration to eliminate cumulative sensor errors still relies on GNSS signal measurements. This implies a persistent error source attributed to MP effects.

With the increasing demand for measurement accuracy, integrity and low power consumption, coupled with the rising availability of pilot signals in the GNSS signal landscape, the principles of long CI find their application in enhancing sensitivity and mitigating MPs (Pany et al. (2009); Cheng et al. (2018); Groves et al. (2020); Xie (2023)). Long CI, when combined with vector tracking in a subset of the position parameters such as the velocity domain (Cheng (2016)), or in the clock offset/drift domain (Cheng (2017)), allows for assessing the log-likelihood metrics in a fine grid with relatively low complexity. In mobile applications, a moving RX with SupercorrelationTM (Faragher et al. (2018)) samples the space during a long CI, which can be interpreted as beam-forming performed by a virtual antenna array. In this regard, a beam towards any arbitrary direction can be formed. Coupled with a suitable counterweight, this can be used for the anti-jamming/spoofing purpose (Faragher et al. (2019)). This way, the CI gain manifests itself in a spatial selectivity, yet only in the presence of antenna motion. The primary challenge in long CI is the signal frequency variation caused by three main factors: the Doppler effect from satellite movement, local clock instability, and the Doppler effect from RX's and/or reflector's mobility. While Doppler dynamics due to satellite movement can be relatively easily predicted and tracked, the other two factors usually require estimation. Without an IMU, dynamics caused by RX motion can be estimated by using alternative transforms that address frequency rates. One of such methods is the fractional Fourier transform, which utilizes a linear chirp basis, as proposed by Luo et al. (2023).

The novelty and contribution of this paper can be summarized as follows.

- We propose a LOS detection/estimation scheme based on SALI, eliminating the need for tracking loops, and reducing the correlation among consecutive measurements. The key departure from the parametric techniques based on Expectation-Maximization (EM) (Feder and Weinstein (1988); Van Nee (1992); Fleury et al. (1999); Lentmaier et al. (2008)) lies in the reduction of the EM iterations to just once, exclusively in the LOS parameter subspace. After a single correlation, realized by SALI using the IMU-aided LOS estimates as initial values, a part of the MPs is already suppressed by the long CI due to significantly mismatched MP signatures to the LOS hypothesis, and the rest MPs remain detectable. In the single M-step, MPs are detected when possible in addition to the LOS, to distinguish the LOS. Moreover, we introduce methods to exclude NLOS measurements, and indicate the level of MP interference that potentially affects a measurement. We analyze the reference time of the input signal, to which a SALI-based LOS measurement is applicable, and the robustness of the PE/ME close-loop, in the presence of measurement and PE feedback errors.
- In the frequency domain, we achieve MP separation and suppression by SALI, particularly for time-variant channels with non-negligible Doppler spread. In the delay domain, we apply linear interpolators to the correlator taps, aiming to 1) mitigate the remaining MPs with indiscernible Doppler to the LOS, in one step without DLL iterations, and 2) enhance the precision of the coarse provisional delay estimates on a discrete grid. Specifically, we introduce two interpolators that mimic the behavior of the well-known DD-DLL, while approaching the set-point in one step, taking advantage of the increased Signal-to-Noise Ratio (SNR) at all correlator taps due to SALI.
- We propose utilizing the Gabor transform with an adaptive Gaussian window curvature for the long CI. This approach complements the Discrete Fourier Transform (DFT), targeting not only the suppression of spectral side lobes, but also the enhancement of robustness against residual Doppler dynamics due to sensor or heading errors.
- In the absence of the IMU-aided motion compensation, we cut the CI length from 1s to 100ms, without explicitly estimating the LOS Doppler rate in each measurement epoch. Here, we use a Kalman Filter (KF) to exploit the channel dynamics across consecutive snapshots.
- Our study showcases the feasibility of SALI as a low-complexity scheme running on a single-chip GNSS RX, with identical hardware accelerators serving a conventional SF RX. In our design, both the Measurement Engine (ME) and the PE efficiently share the computational capacity of a single ARM Cortex M3 processor in real-time.
- For a fair comparison, the SALI-based SF RX is tested against the same SF architecture employing conventional carrier and code tracking loops. In each individual environment, including open sky, foliage, suburban and deep urban, from a

comprehensive road test campaign across the globe, we compare the delay and Doppler error statistics, as well as position fix statistics, in terms of accuracy, integrity and availability w.r.t. the required 2m accuracy using L1 signals. In addition, we compare and interpret those statistics for different C/N0 as well as for different rover speeds.

The remainder of the paper is organized as follows. We first present a signal model that combines LOS geometry and stochastic properties of time-variant MP channels. Next, we introduce SALI and discuss its sensitivity to various MP effects and sensor-aiding errors. Based on the SALI output, we derive the LOS detection and estimation algorithm, and discuss the reference time of the measurements, robustness and the MP mitigation capability in the presence of various channel characteristics. We then introduce three low-complexity delay interpolators, and discuss their performance w.r.t. the well-known DD-DLL. To further enhance the robustness in the absence of sensor aiding, we propose to use shorter CI time in combination with a KF. We present road test results related to measurement errors, positioning accuracy, integrity, and availability. These results are categorized by various environments, and we discuss their implications. Finally, we summarize the results, and present the advantage and disadvantage of the proposed SALI-based SF RX, along with an outlook for the future research.

SIGNAL MODEL

Throughout the paper, we use the following mathematical notations. A function of a continuous or discrete variable u is denoted by $y(u)$ or $y[u]$, respectively. Parameters can be added to a function, such as ψ and ρ in $y(u; \psi, \rho)$. The first derivative, absolute value and complex conjugate of y are denoted y' , $|y|$ and y^* , respectively. The imaginary unit $\sqrt{-1}$ is denoted by j . The Euclidean distance between two functions $y(\cdot)$ and $z(\cdot)$ is denoted by $\|y(\cdot) - z(\cdot)\|$. A Random Variable (RV) without any accent denotes the unknown truth, $\hat{\cdot}$ denotes the estimate, and $\bar{\cdot}$ denotes the corresponding estimation error, respectively. Since the latter is also an RV, its value can be estimated, and the error estimate is denoted by $\hat{\hat{\cdot}}$ consequently. $\mathbb{P}(\cdot)$ denotes the probability of an event.

The baseband received signal, $r(t)$, is modeled as a LOS component indexed by $l = 0$ superimposed with L MP components indexed by $l = 1, 2, \dots, L$ in (1). The sum is corrupted by Additive White Gaussian Noise (AWGN) $n(t)$ with a double-sided noise power spectral density of $N_0/2$.

$$r(t) = \sum_{l=0}^L \sum_{k=0}^{K-1} h_l[k] e^{j\phi_l[k]} x(t - kT - \tau_l[k]) + n(t). \quad (1)$$

Here, the piece-wise constant complex amplitude, $h_l[k]$, for path l in code period k during interval $t \in [kT, (k+1)T)$ is modelled as an RV following certain fading statistics. T represents the spreading code period, such as $T = 1\text{ms}$ for GPS L1CA. The geometry-based carrier phase and propagation delay for path l in code period k are represented by $\phi_l[k]$ and $\tau_l[k]$, respectively. Both are RVs, influenced by the RX, satellite movement and the ray-tracing geometry of path l . The spreading code waveform denoted by $x(t)$ is defined within the interval of $t \in [0, T)$, as a result of a chip pulse waveform convolved with the binary spreading code sequence. This implies that code period $k = 0$ begins with $t = 0$. For simplicity, we also assume that the modulating bits or secondary codes are removed during the entire observation interval of $t \in [0, KT)$, which is $\sim 1\text{s}$ in our study. An obstructed path corresponds to $|h_l[k]| = 0$. The magnitude and phase of $h_l[k]$ undergo a relatively small variation in the presence of Rician fading. This can be a visible LOS path or specular reflections with long-lasting lifespans. In case of Rayleigh fading such as in NLOS environments, the phase can be assumed uniformly distributed in $[0, 2\pi)$ and the fading rate is subject to the path coherence time. Here, the composite carrier phase in $h_l[k] e^{j\phi_l[k]}$ can be dominated by the geometric or stochastic characteristics. Furthermore, for terrestrial mobile channels, we can approximate the geometric LOS dynamics at $t = kT$ with instantaneous velocities, $v_0[k]$, which results from projecting the relative velocity between the RX and the satellite to the LOS vector. The LOS range in code period k differs from that in $k = 0$ by $d_0[k] = \sum_{i=0}^{k-1} v_0[i]T$. This leads to the following LOS carrier phase, delay and frequency in code period k :

$$\phi_0[k] = \phi_0[0] + \frac{2\pi}{\lambda} d_0[k], \quad \tau_0[k] = \tau_0[0] + \frac{d_0[k]}{c}, \quad \text{and} \quad f_0[k] = \frac{v_0[k]}{\lambda},$$

with λ and c being the wavelength and the speed of light. For MPs ($l > 0$), the same mechanics govern the evolution of $\phi_l[k]$, $\tau_l[k]$ and $f_l[k]$. However, the propagation path geometry is generally unknown without information about reflector positions, and projecting the RX velocity to ray-tracing paths is not feasible. On the other hand, it is not required to do so because the long CI is exclusively computed w.r.t. the LOS component, as will be elaborated in the next sections.

Note that in (1), we disregard the impact of local clock instability, which can lead to Doppler and delay dynamics in the received signal $r(t)$ as well. The rationale behind this simplification is the fact that our hardware platform, which effectively manages thermal coupling, exhibits satisfactory frequency stability for the observation window of 1s, even in the presence of vibration.

This assumption is not overly ambitious for dedicated GNSS RXs, and we can tolerate minimal clock drift and offset in each measurement epoch without the need to estimate and remove them separately.

SALI

We follow a feedforward detection and estimation strategy, under the maximum likelihood framework. To that end, we model the local replica in (2), focusing exclusively on the LOS component with the frequency and delay hypothesis, f, τ .

$$s_0(t; f, \tau) = \sum_{k=0}^{K-1} e^{j(\hat{\phi}_0[k] + 2\pi f k T)} x(t - kT - (\hat{\tau}_0[k] + \tau)). \quad (2)$$

Here, $\hat{\phi}_0[k]$ and $\hat{\tau}_0[k]$ are the IMU-aided initial phase and delay estimates of the LOS path. They are used to control the carrier and code generator, respectively, to track the signal in code period $k = 0, 1, \dots, K - 1$. In an SF RX, both are derived from the IMU-aided LOS velocity estimate, $\hat{v}_0[k]$, as follows:

$$\hat{\phi}_0[k] = \hat{\phi}_0[0] + \frac{2\pi}{\lambda} \sum_{i=0}^{k-1} \hat{v}_0[i]T, \quad \hat{\tau}_0[k] = \hat{\tau}_0[0] + \frac{1}{c} \sum_{i=0}^{k-1} \hat{v}_0[i]T,$$

where $\hat{\phi}_0[0]$ and $\hat{\tau}_0[0]$ represent the initial phase and delay estimate that are applied to the carrier and code generator at $t = 0$. In addition, the IMU-aided LOS frequency estimate during $t \in [kT, (k+1)T)$ is given by $\hat{f}_0[k] = \hat{v}_0[k]/\lambda$. Without compromising the system behavior, we ignore the latency of the feedback control in this section. Upon defining the phase and delay difference of all $L + 1$ paths against the LOS estimates, as $\bar{\phi}_l[k] = \phi_l[k] - \hat{\phi}_0[k]$ and $\bar{\tau}_l[k] = \tau_l[k] - \hat{\tau}_0[k]$, the correlation function of $s_0(t; f, \tau)$ and $r(t)$ reads

$$\begin{aligned} R(f, \tau) &= \int_{-\infty}^{\infty} W(t)r(t)s_0^*(t; f, \tau)dt \\ &\approx \nu + \sum_{l=0}^L \sum_{k=0}^{K-1} W[k]h_l[k]e^{j(\bar{\phi}_l[k] - 2\pi f k T)} \int_{-\infty}^{\infty} x(t - kT - \tau_l[k])x^*(t - kT - (\hat{\tau}_0[k] + \tau))dt \\ &= \nu + \sum_{l=0}^L \sum_{k=0}^{K-1} W[k]h_l[k]e^{j\bar{\phi}_l[k]} \Lambda(\bar{\tau}_l[k] - \tau) e^{-j2\pi f k T} \\ &= \nu + \sum_{l=0}^L R_l(f, \tau). \end{aligned} \quad (3)$$

Here, RV $\nu = \int_{-\infty}^{\infty} n(t)s_0^*(t; f, \tau)dt$ with variance $KTN_0/2$ represents the integrated AWGN, and a Gaussian window, $W(t) = \exp[-\pi\alpha(t - KT/2)^2]$ sampled at $t = kT$ to $W[k]$, is applied to the code correlation samples. Considering the MP signal model (1) and the LOS hypothesis (2) without any MPs, we recognize that the correlation in (3) corresponds to the evaluation of a partial log-likelihood function. In the context of the Space-Alternating Generalized EM (SAGE) algorithm (Fleury et al. (1999)), (3) prepares for the M-step of the first EM iteration in the LOS parameter subspace, using the sensor-aided LOS estimates as initial values. The sum in (3) includes the auto-correlation of the LOS path, plus all the cross-correlations between the LOS hypothesis and MPs. This ensures a maximum CI gain for the LOS path, while allowing MPs to either concentrate in their respective frequency bins, or diminish due to the cross-correlation with significantly mismatched signatures compared to the LOS hypothesis. In (3), The correlation of the LOS hypothesis and path $l \geq 0$ is represented by

$$R_l(f, \tau) = \sum_{k=0}^{K-1} W[k]h_l[k]e^{j\bar{\phi}_l[k]} \Lambda(\bar{\tau}_l[k] - \tau) e^{-j2\pi f k T} \quad \text{with} \quad \Lambda(\tau) = \begin{cases} 1 - |\tau|/T_c & \forall \tau \in [-T_c, T_c], \\ 0 & \text{elsewhere.} \end{cases} \quad (4)$$

In this context, Gaussian windowing followed by the DFT effectively computes the Gabor transform of the code correlation samples, considering a window displacement fixed to the center of the observation interval KT . A tuning coefficient α shapes the curvature of the window. Moreover, we assume infinite bandwidth applied to a train of rectangular chipping pulses with a chip duration of T_c . This leads to an ideal spreading code auto-correlation function, $\Lambda(\tau)$. For an observation window of $\sim 1s$, we approximate the velocity and range difference of path l in code period k relative to the IMU-aided LOS estimates using a

second-order dynamics model as follows:

$$\bar{v}_l[k] = v_l[k] - \hat{v}_0[k] \approx \bar{v}_l[0] + \bar{a}_l kT, \quad \bar{d}_l[k] = d_l[k] - \hat{d}_0[k] \approx \bar{d}_l[0] + \bar{v}_l[0]kT + \frac{1}{2}\bar{a}_l k^2 T^2. \quad (5)$$

Here, $\bar{v}_l[0]$ and $\bar{d}_l[0]$ are the initial velocity and range difference between path l and the LOS estimate. A constant rate of change in velocity difference is denoted by \bar{a}_l . Note that $\bar{v}_l[k]$ and $\bar{d}_l[k]$ exclusively represent estimation errors of the LOS path for $l = 0$, and for any other paths they are merely a difference in the cross-correlation. However, for simplicity, we refer to them and the same type of differences collectively as “error” for all paths $l \geq 0$ in the following discussion. By considering the wavelength and the speed of light, we can convert (5) to the following phase, delay and frequency error of path l , respectively,

$$\bar{\phi}_l[k] = 2\pi \frac{\bar{d}_l[k]}{\lambda} = 2\pi \frac{\bar{d}_l[0]}{\lambda} + 2\pi \frac{\bar{v}_l[0]}{\lambda} kT + \frac{1}{2} \frac{2\pi \bar{a}_l}{\lambda} k^2 T^2 = \bar{\phi}_l[0] + 2\pi \bar{f}_l[0] kT + \pi \bar{f}'_l k^2 T^2, \quad (6)$$

$$\bar{\tau}_l[k] = \frac{\bar{d}_l[k]}{c} = \frac{1}{c} \left(\bar{d}_l[0] + \bar{v}_l[0]kT + \frac{1}{2}\bar{a}_l k^2 T^2 \right), \quad (7)$$

$$\bar{f}_l[k] = f_l[k] - \hat{f}_0[k] = \frac{\bar{v}_l[k]}{\lambda},$$

with $\bar{\phi}_l[0] = 2\pi \bar{d}_l[0]/\lambda$, $\bar{f}_l[0] = \bar{v}_l[0]/\lambda$ and $\bar{f}'_l = \bar{a}_l/\lambda$ representing the initial phase error, initial frequency error and constant frequency error rate of path l , respectively. Notice that the underlying error signal exhibits carrier dynamics of a linear chirp.

FIGURE 1 shows the simulation of an artificial scenario with a LOS path plus one MP, both originating from a satellite at a low elevation of $\sim 0^\circ$. Here, a rover initially travels eastward in a straight line at 50 km/h for 1s. In the following second, it takes a 90° turn to the south along a circular path while maintaining the constant linear speed, reaching an acceleration of 2.2g. The LOS originates from the east, and the MP from azimuth 210° persists for the entire 2s. A heading error of constant 3° is assumed for the IMU-aided velocity estimate throughout the 2s duration. Doppler values in FIGURE 1b reveal that both the frequency error and its rate remain consistently low for the LOS signal, irrespective of linear motion or navigating a curve because the LOS path is closely tracked by the IMU-aided SF solution. However, for MP components, the frequency error can be large already during the linear motion. Additionally, the frequency error rate can be high given non-zero rover acceleration.

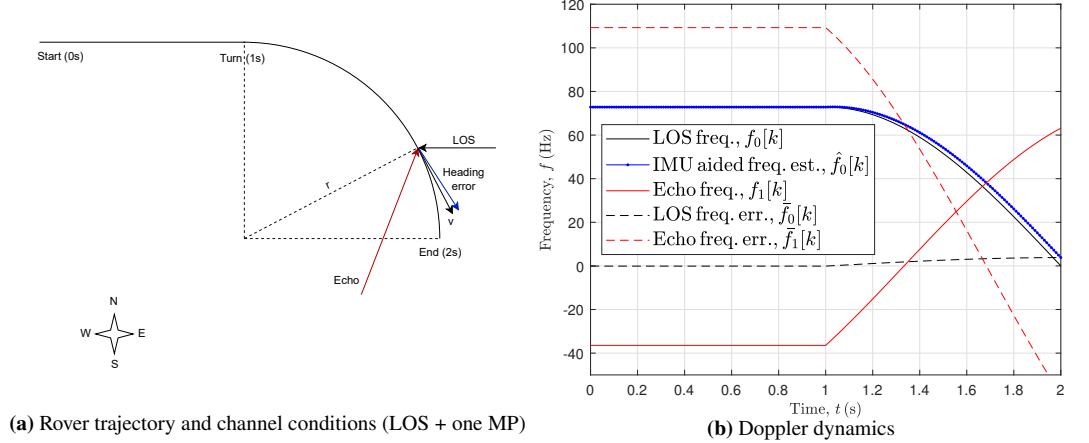


Figure 1: IMU-aided LOS tracking in MP environment

The impact of $\bar{f}_l[0]$, \bar{f}'_l and $\bar{\tau}_l[k]$ on the correlation in (4), given small frequency errors, can be analyzed by neglecting the amplitude variation in $h_l[k]$, and approximating the vicinity of the correlation peak by splitting the spreading code correlation, $P_l(\tau)$, from the sinusoidal correlation, $Q_l(f)$, which is the Gabor transform of a linear chirp, as follows:

$$R_l(f, \tau) \approx P_l(\tau)Q_l(f) \quad \text{with} \quad P_l(\tau) = \frac{h_l[0]}{K} \sum_{k=0}^{K-1} \Lambda(\bar{\tau}_l[k] - \tau) \quad \text{and} \quad Q_l(f) = \sum_{k=0}^{K-1} W[k] e^{j\bar{\phi}_l[k]} e^{-j2\pi f kT}. \quad (8)$$

Both $P_l(\tau)$ and $Q_l(f)$ have closed-form expressions. However, it suffices to analyze their global maximum, which is crucial for subsequent path detection and estimation. According to (7), $\bar{\tau}_l[k]$ wanders along $\Lambda(\cdot)$ for non-zero $\bar{f}_l[k]$ when summing over $k \in [0, K-1]$ for any specific delay hypothesis τ . As a result, the peak of $P_l(\tau)$ is flattened and displaced in comparison to the case of $\bar{f}_l[k] = 0$. It can be shown, that for constant frequency error with $\bar{f}'_l = 0$, $P_l(\tau)$ peak is deformed to a parabola, which is displaced by $(KT/2)\bar{f}_l[0]\lambda/c$, as shown in FIGURE 2a. The resulting delay associated with the maximum corresponds to the delay of the signal at $t = KT/2$. If, however, the frequency error is time-variant with an unknown rate of change, as we

intentionally do not estimate \hat{f}'_l for the sake of computational complexity, the offset of the $P_l(\tau)$ peak remains uncertain in general. This implies a minimal yet noticeable delay estimation error, especially in high dynamic situations, and even in perfect open sky conditions.

On the other hand, even without the knowledge of \hat{f}'_l , the Gabor transform consistently shifts the spectrum peak of $Q_l(f)$ by $(KT/2)\hat{f}'_l$ for small \hat{f}'_l , as shown in FIGURE 2b. Here, the frequency error estimate associated with the spectrum peak is relevant to the time at $t = KT/2$, regardless of the value of \hat{f}'_l . With an increasing frequency error rate, the spectrum becomes more dispersive, resulting in a lower global maximum that can even appear at a different frequency than $(KT/2)\hat{f}'_l$, as shown in FIGURE 2b. This phenomenon is counterproductive for the LOS Doppler estimation.

To mitigate its impact, one approach is to use a narrower window by choosing a larger value for α . While the narrower window resists larger Doppler error rates and attenuates more spectrum sidelobes, it reduces the peak SNR, diminishing the detection sensitivity. In practice, this leads to a trade-off conditioned by propagation channel characteristics and the sensor aiding error properties.

Obviously, the Doppler error rate over the integration interval is a critical factor that can impact the robustness of SALI. Simulations of the IMU-aided frequency tracking in our study exhibit a maximum of 4Hz/s under a relaxed heading error assumption. This justifies the omission of explicit frequency error rate estimation. Considering MP components, a large frequency error rate can be rather beneficial for LOS estimation, especially when the initial MP Doppler is hardly discernible from that of the LOS. In such cases, the increased frequency error rate scatters the MP power, effectively reducing its interference to the LOS correlation.

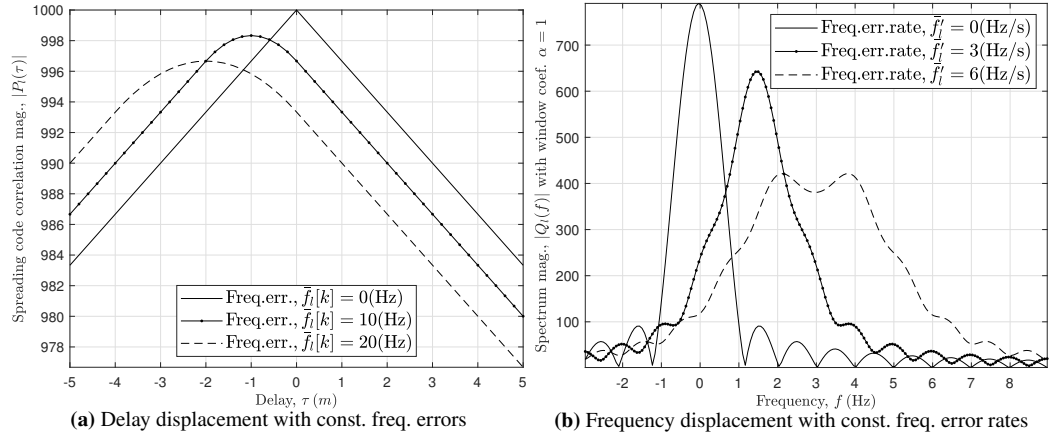


Figure 2: Impact of frequency error and its rate on 1s CI

SALI-BASED LOS DETECTION AND ESTIMATION

The problem of LOS parameter measurement can be formulated as deriving an estimate, $\hat{f}_0, \hat{\tau}_0$, of the frequency and delay errors, $\hat{f}_0[k], \hat{\tau}_0[k]$, associated with the initial sensor-aided estimates, $\hat{f}_0[k], \hat{\tau}_0[k]$, for $k = \lfloor K/2 \rfloor$, by evaluating $R(f, \tau)$. In the SAGE context, as explained w.r.t. the correlation in (3), the procedure here corresponds to the single M-step in the LOS parameter subspace. Since no further EM iterations are planned to estimate each path individually, all detectable paths after the long CI must be identified during the single M-step to differentiate the LOS. In particular, we apply a detection threshold η to $|R(f, \tau)|$ to detect multiple local maxima associated with individual paths. This resembles an aided acquisition, with the key difference being that multiple paths, rather than just the strongest one, are detected. From the detected paths, we identify the one with the shortest delay as the LOS. Its discrete delay and frequency values linked to the correlation peak on the hypothesis grid are then interpolated to obtain a more precise estimate. Finally, we conclude the measurement epoch by reporting the LOS estimates w.r.t. the reference time of $t = KT/2$. FIGURE 3 presents a block diagram that can implement ALGORITHM 1 in a single-chip SF RX with conventional GNSS accelerators.

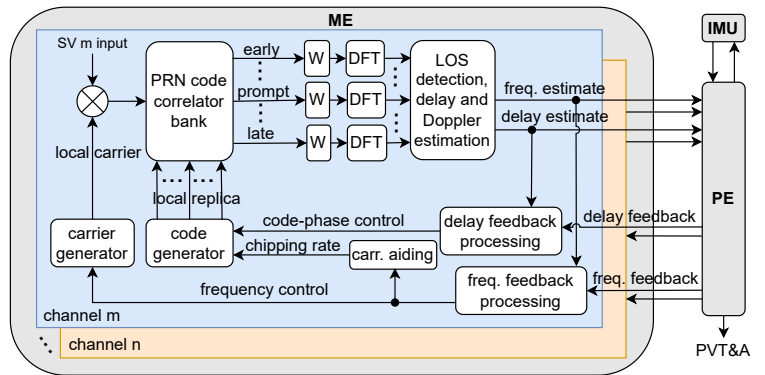


Figure 3: SALI-based LOS detection and estimation

Algorithm 1 SALI-based LOS detection and estimation

Require: Real-time tracking of $r(t)$ with ME controlled by $\hat{\phi}_0[0], \hat{\tau}_0[0], \hat{f}_0[k]$ for $k = 0, 1, 2, \dots, K - 1$;
Ensure: Approximation (5) holds, and \hat{f}'_0 is sufficiently small;
Require: Definition of hypothesis windows, $\tau \in [\tau_{\min}, \tau_{\max}]$ and $f \in [f_{\min}, f_{\max}]$;
Ensure: Proper hypothesis window coverage: $|\hat{\tau}_0[k] - \tau| \leq T_c$ and $f_{\min} \leq f_0[0] + \hat{f}'_0 KT/2 \leq f_{\max}$;

- 1: $\hat{\tau}_0 \leftarrow \infty$;
- 2: $\hat{f}_0 \leftarrow 0$;
- 3: **for all** $f \in [f_{\min}, f_{\max}]$ **do**
- 4: $\tau_l \leftarrow \arg \max_{\tau} |R(f, \tau)|$, c.f. (3);
- 5: **if** $|R(f, \tau_l)| > \eta$ **then**
- 6: $\tau_l \leftarrow$ interpolate τ in the vicinity of τ_l based on $|R(f, \tau)|$ given specific f ;
- 7: **if** $\tau_l < \hat{\tau}_0$ **then**
- 8: $\hat{\tau}_0 \leftarrow \tau_l$;
- 9: $\hat{f}_0 \leftarrow f$;
- 10: **end if**
- 11: **end if**
- 12: **end for**
- 13: $\xi \leftarrow$ calculate MP indicator $\|R(\hat{f}_0, \tau) - \Lambda(\tau - \hat{\tau}_0)\|$;
- 14: $\hat{f}_0 \leftarrow$ interpolate f in the vicinity of \hat{f}_0 based on $|R(f, \hat{\tau}_0)|$ given $\hat{\tau}_0$;
- 15: **return** LOS delay estimate, $\hat{\tau}_0[\lfloor K/2 \rfloor] + \hat{\tau}_0$, LOS frequency estimate, $\hat{f}_0[\lfloor K/2 \rfloor] + \hat{f}_0$ and MP indicator, ξ ;

ROBUSTNESS

In general, SALI operates with minimal reliance on prior knowledge or assumptions about the channel model, e.g., the number of existing paths. Despite similarities in the first EM-iteration w.r.t. the LOS, the weak dependency on channel assumptions distinguishes SALI from the iterative joint parameter estimators such as SAGE, and contributes to its robustness in single measurement epochs. However, since we are dealing with a feedback control loop involving both ME and PE, robustness over consecutive epochs regarding initial conditions in each epoch is still a critical concern. Beyond the earlier discussion on frequency error rates, it's crucial to examine the frequency and delay hypothesis window. This window must strike a balance: it needs to be large enough to encompass the LOS frequency and delay errors associated with IMU-aided signal tracking over the integration interval, while also being as small as possible to achieve two objectives: 1) keeping as many as possible MP correlation peaks outside the LOS hypothesis window to guard against the NLOS situation, and 2) ensuring a high resolution of the hypothesis grid, given limited computational resources. Furthermore, it's essential to take into account additional errors arising from feedback control latency when determining the size of the hypothesis window, while minimizing the latency remains best practice. Neglecting the impairment from remaining MP interference, interpolation error and the unknown frequency error rates, the estimator can provide unbiased estimates of LOS parameters. This process augments the initial IMU-aided estimates by leveraging the received signal within the maximum likelihood framework. Error propagation between ME and PE is inherently precluded by design.

If the initial LOS estimate from PE is biased in such a way that the true LOS parameters lie outside the hypothesis window, we're confronted with a false NLOS situation, where all correlation peaks within the window are attributed to MPs. Whether in a false or true NLOS scenario, our preference in the context of safety critical applications is to refrain from producing a measurement and default to an NLOS declaration, awaiting the next measurement epoch. In addition to SALI's inherent capability of excluding MP correlation peaks out of the narrow hypothesis window, various methods can be employed to enhance the NLOS exclusion. One possibility is to set the detection threshold η to a C/N0 level, below which the majority of the detected paths are expected to originate from NLOS Rayleigh fading sources according to channel models such as Lehner (2007). The other approach stems from the SNR saturation effect observed in long CI under Rayleigh fading conditions (Broumandan et al. (2011)). Specifically, the resulting SNR gap between a Rician and Rayleigh faded signal grows with the increasing CI length and the decreasing channel coherence time. Correspondingly, we can monitor the rate of SNR growth during each SALI epoch to differentiate the realization of a Rician (LOS present) and Rayleigh faded NLOS scenario.

Another robustness concern is the outage of IMU aiding, which is not only an abnormal state but also expected in practice, such as before sensor calibration is completed. This typically prevents the long CI due to the lack of motion compensation required by SALI. As a backup solution, we propose to reduce the CI length from 1s to 100ms without estimating the Doppler rate in each measurement epoch. The resulting spectrum dispersion of the LOS component remains acceptable up to 2g acceleration. The LOS parameter estimates, obtained from 100ms snapshots, drive a KF in individual tracking channels. This KF tracks the inherent motion dynamics, in conjunction with a 1Hz navigation rate. Note that in the nominal operation mode, as previously

explained, the exchange of information between the ME and the PE occurs sequentially, making the additional tracking loops in ME redundant, as the navigation filter effectively handles rover dynamics tracking. Although it's possible to configure the nominal mode for 100ms CI, which implies a navigation rate of 10Hz, the MP mitigation effect will be diminished because of the tenfold degraded frequency resolution compared to the 1s SALI. This also implies that for high-rate position output, it's advisable to run sensor-propagation between two consecutive nominal SALI-based solutions with a 1s CI.

DELAY AND FREQUENCY INTERPOLATION

In practice, the correlation function $R(f, \tau)$ is sampled on a discrete grid over f and τ . To achieve higher accuracy, an interpolation is needed. In the frequency domain, a Quadratic Interpolation (QI) of frequency bins around the spectrum peak already provides sufficient accuracy in STEP 14 of ALGORITHM 1. In scenarios with minimal Doppler spread, SALI's capability of separating MPs by frequencies diminishes. To mitigate MPs, which produce their spreading code correlation function $P_l(\tau)$ on the same frequency as the LOS, we apply the design principles of multiple early-late gates with narrow tap spacing in a one-shot delay interpolator without DLL iterations, as shown in STEP 6 of ALGORITHM 1. All three delay interpolators below operate with identical prerequisites, namely the maximum of $|R(f, \tau)|$ for a specific f has been detected as $A[m] := |R(f, \tau[m])|$ at delay tap m . Its relative delay to the reference delay, $\hat{\tau}_0[[K/2]]$, configured at tap 0 is $m\Delta$ with Δ denoting the tap spacing. Additional early-late taps, $A[m \pm 1], A[m \pm 2], \dots, A[m \pm M]$, are also required for the same f . As a benchmark to compare the interpolator performance, we consider the LOS delay estimation error in a static two-path model in FIGURE 4, where the steady state error of the DD-DLL in FIGURE 4a serves as the reference.

- Double-Delta Interpolation (DDI) using seven taps: For sufficiently large M , we can fit the output of a DD discriminator associated with tap i ,

$$S[i] = 2(A[i - 1] - A[i + 1]) - (A[i - 2] - A[i + 2]),$$

to the delay continuous S-curve, and solve the zero-crossing equation for the refined delay estimate. As an example, we sample the S-curve of a DD discriminator with $\Delta = T_c/16$ at three points, $i = m, m \pm 1$, i.e., seven taps are involved. FIGURE 4b shows the resulting delay error given tap 0 is biased by $1.74 \cdot \Delta$, for instance.

- Early-Late Line Intersection (ELLI) using five taps: The interpolated delay estimate is obtained by identifying the intersection point between the line defined by the early taps ($A[m - 1]$ and $A[m - 2]$) and the line defined by the late taps ($A[m + 1]$ and $A[m + 2]$) as follows:

$$\hat{\tau}_0 = \Delta \left(m + S[m](A[m - 2] + A[m + 2] - A[m - 1] - A[m + 1])^{-1} \right), \quad (9)$$

where the numerator, $S[m]$, is the DD discriminator output associated with the correlation peak at tap m . At the DLL set-point, $S[m]$ approaches 0, which implies that both lines would intersect exactly at tap m in the middle of tap $m - 1$ and $m + 1$. Therefore, ELLI using five taps in (9) mimics the behavior of the DD-DLL in one shot. As long as tap m is close to DD-DLL's set-point, a similar performance can be expected. FIGURE 4c shows the resulting delay errors, which can overreact to the antipodal MP phase, compared to the DD-DLL.

- QI of three taps, $A[m - 1], A[m]$ and $A[m + 1]$, can be employed if the number of taps are limited (FIGURE 4d).

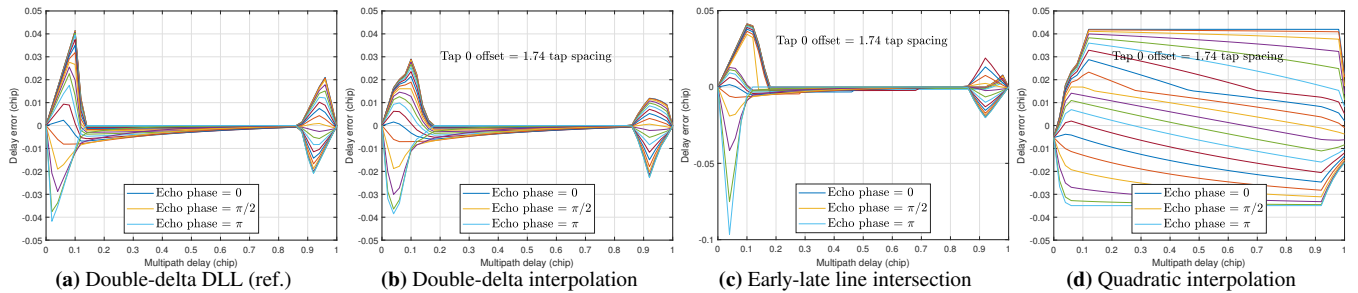


Figure 4: LOS + one MP with -3dB, delay $\in [0, T_c]$, phase $\in [0, \pi]$, and identical frequency w.r.t. LOS, tap spacing = $T_c/16$

In FIGURE 5, we average the error over all MP delay and phase values to obtain a mean MP error for any given tap 0 bias. It becomes apparent that DDI and ELLI exhibit slight outperformance and underperformance compared to DD-DLL on average, while both are more MP resistant than the simple QI. However, unlike conventional DLLs, a one-shot estimator on a discrete grid followed by an interpolator is inconsistent in the sense, that the final estimate depends on the initial timing condition, as shown by the Δ -periodic pattern in the mean errors in FIGURE 5. The interpolation errors are mostly negligible w.r.t. the MP errors, yet they persist with the increasing amount of observation samples. The primary benefit of a one-shot interpolator over a DLL lies in its superior code tracking agility, reducing the correlation among consecutive measurements.

Despite all the MP mitigation capabilities and NLOS exclusion methods enabled by SALI, there remains a rest of MP interference, particularly in NLOS scenarios with low Doppler dispersion. To address this, we calculate an MP indicator ξ in each measurement, as the Euclidean distance between the expected spreading code correlation function and the actual obtained correlation values, as shown in STEP 13 of ALGORITHM 1. This enables PE to properly assign weights to individual measurements in the navigation filter.

MULTIPATH MITIGATION

The MP mitigation capability of SALI can be discussed w.r.t. stochastic characteristics of mobile channels. In particular, the coherence time, Doppler and delay spread play an important role.

The LOS path may be attenuated when it penetrates a medium. Simultaneously, it can be superimposed with MPs from local scatterers. When the delay of the latter is indiscernible w.r.t. the LOS for the given signal/RX bandwidth, a Rician faded amplitude, $h_0[k]$, arises. A high K-factor corresponds to a dominant LOS component with long coherence distance in general. On the MP side, channel models typically prescribe a Rayleigh fading process with a short coherence distance of several wavelengths, statistically. In real world, however, it's not uncommon to observe specular reflections with long coherence distance that only appears as tail probabilities in channel models. For instance, consider a scenario where a RX is moving in a straight urban canyon with sky-scrapers at both ends. The few visible satellites with low elevation are located either in the front or back of the RX. They may be subject to specular reflections with a relatively long coherence distance. If LOS is obstructed, we may only observe MPs, which can be either Rayleigh faded due to diffuse reflections or Rician faded due to specular reflections.

On the RX side, SALI with a CI length of ~ 1 s can achieve a frequency resolution of about 1Hz, which is $50\times$ narrower than that of the 20ms CI. FIGURE 6a and 6b illustrate the impact with a simulation of a LOS path plus two specular MPs. All three paths are close to each other in a frequency range of several Hz, and the LOS component is 3dB weaker than either of the MPs. The frequency error rates of the three paths are $\bar{f}'_0 = \bar{f}'_1 = 0$ and $\bar{f}'_2 = 10\text{Hz/s}$, respectively. The latter spreads MP2's spectrum $Q_2(f)$ and suppresses its peak. As a result, only correlation peaks of the LOS and MP1 are detectable after 1s CI, although both are merely 2.5Hz apart, as shown in FIGURE 6b. Furthermore, the resulting correlation function after 1s CI encompasses more likely a single $\Lambda(\cdot)$ for each individual path frequency. This facilitates accurate delay estimation, especially when the RX bandwidth is sufficiently large. In contrast, the 20ms CI can hardly separate the three paths with the frequency resolution of 50Hz, and the stronger MPs can dominate the delay estimation in FIGURE 6a.

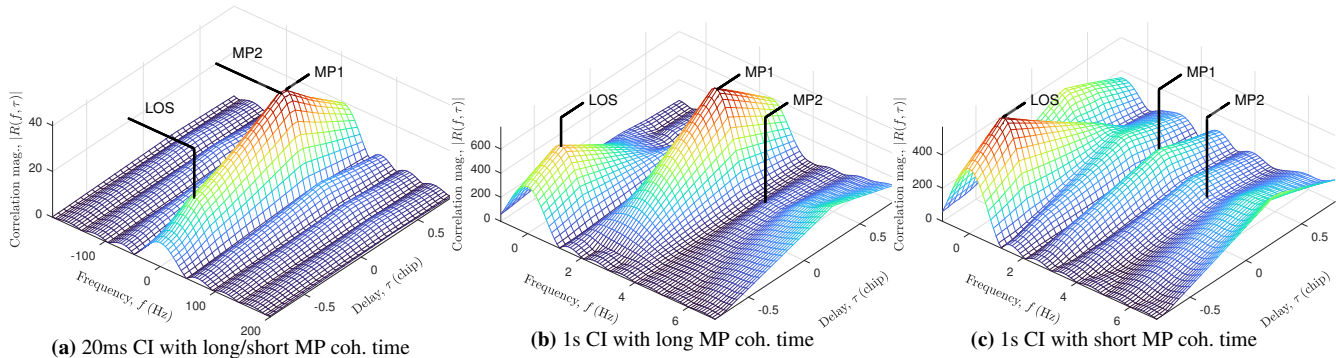


Figure 6: Correlation in a three-path model with different freq. error rates and MP coh. time

For moving RXs, Doppler spread is inversely proportional to the coherence time, which is directly proportional to the coherence distance, and inversely proportional to the rover speed. Based on the same model as in FIGURE 6b, if we assume diffuse instead of specular reflections for MP1, short coherence time arises. Assuming a rover speed of 50km/h and a coherence distance of $2.5\lambda = 0.5\text{m}$ for L1 carrier frequencies (ITU-R P.681-8 (2015)), we obtain a coherence time of about 36ms. During a 1s CI, the RX should experience roughly 28 uncorrelated realizations of the Rayleigh faded complex amplitude $h_l[k]$ with uniformly distributed phase between 0 and 2π . The sum of those in (4) partially cancels out, given indiscernible delays. As a result, MP1's power is suppressed, as shown in FIGURE 6c. Notably, this effect does not occur with the $\leq 20\text{ms}$ CI, because the integration length is too short to experience uncorrelated MPs sequentially.

Integrating beyond the lifespans of MPs is particularly beneficial in urban environments, where the LOS signal is received from one side of the street and successive MPs (with short lifespans) are received from the other side. Such scenarios often arise due to reflections from building fronts along the street when a vehicle passes by at low lateral speeds.

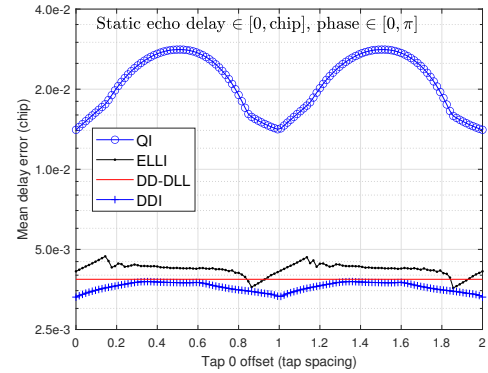


Figure 5: Mean delay errors

The minimal differences in signal power, Doppler, and delay between the LOS and MPs stress delay discriminators utilizing 20ms CI and challenge conventional DLLs, as shown by the distorted peak in FIGURE 7a. In contrast, long CI suppresses uncorrelated MPs, making the LOS correlation peak stand out, and facilitating the delay estimation, as shown in FIGURE 7b. Another similar scenario involves foliage, where the LOS component is attenuated when it penetrates the canopy, yet it remains coherent for a relatively long duration. The leaves and branches within the foliage scatter the impinging waves, resulting in a large amount of MPs with short lifespans.

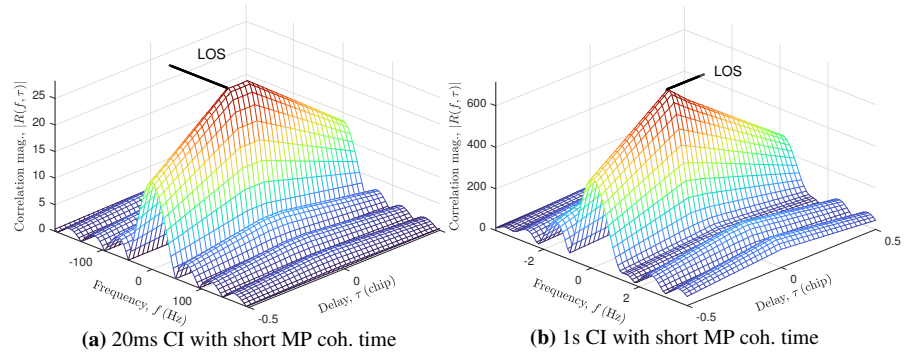


Figure 7: Correlation in a two-path model including near MPs with short coh. time

ROAD TEST CAMPAIGN AND PERFORMANCE ANALYSIS

Extensive test drives are carried out in Thalwil, Zurich, Munich, Hannover, Frankfurt, New York City (NYC), Detroit, Chicago, etc. Test routes are planned carefully in advance to include various environments such as open sky, foliage, suburban, and deep urban areas. A test vehicle is equipped with a GNSS recording box, a consumer-grade IMU and a third-party survey-grade inertial navigation system aided by a high-precision IMU. During test drives, Radio Frequency (RF) signals, the consumer-grade sensor data and vehicle wheel-ticks are recorded synchronously, along with the supposed “truth”, consisting of survey-grade Position, Velocity and Time (PVT) data. In the lab, recorded RF signals together with sensor data are replayed, and fed into identical hardware platforms. The only difference lies in the firmware: one platform runs the SALI-based SF algorithm with ~ 1 s CI, while the other (non-SALI reference) follows the conventional SF implementation with traditional tracking loops based on ≤ 20 ms CI. Both the SALI and the reference RX use the same number of channels to track an identical subset of L1 constellations, exclusively. The performance analysis is conducted in two stages, 1) measurement error analysis followed by 2) position error and integrity analysis. In the first stage, both Pseudo-Range (PR) and Doppler (LOS velocity) errors are analyzed w.r.t. the truth derived from the true PVT data logged. The errors are interpreted w.r.t. different speeds and C/N0 values. In the second stage, horizontal position errors and position accuracy estimates are considered to derive the integrity and availability metrics w.r.t. the required position accuracy of two meters. All the measurement and position errors as well as integrity and availability statistics are analyzed for the four pre-defined environment categories.

MEASUREMENT ERRORS IN ROAD TESTS

Doppler and PR errors of SALI and reference RXs throughout all measurement epochs are grouped into high (>1 m/s) and low rover speed (<1 m/s) categories across all C/N0 values in all road test scenarios. Their Cumulative Distribution Functions (CDFs) are shown in FIGURE 8. In general, Doppler measurements tend to be more accurate at low speeds, as indicated by

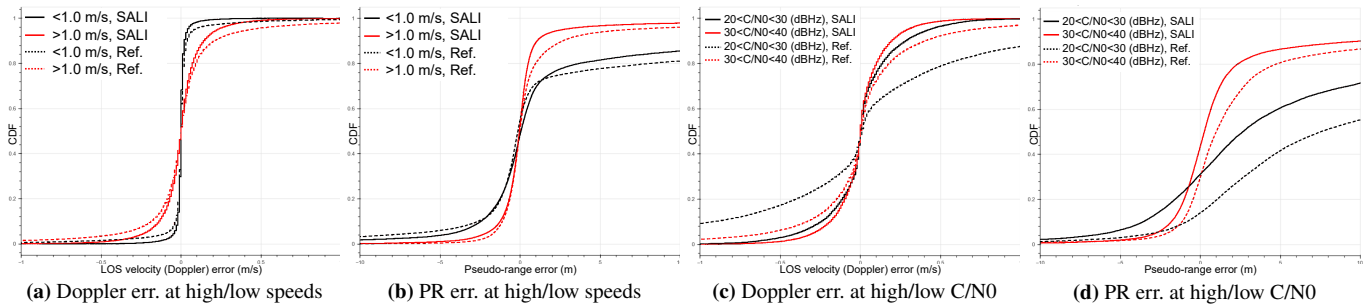


Figure 8: CDFs of Doppler and PR errors at different speeds and C/N0 across all test drives

the sharper CDFs associated with speeds <1 m/s for both SALI and the reference RX in FIGURE 8a. This can be explained by the lower Doppler spread, which accounts for higher consistency in the LOS Doppler, making the frequency estimation more accurate. However, low speeds diminish the DLL filtering effect against time-variant MPs in conventional tracking loops, and reduce SALI’s ability of separating MPs. Consequently, the uncertainty of delays increases in the vicinity of the LOS Doppler, leading to inaccurate delay estimation, as revealed by the flatter CDFs of PR errors associated with speeds <1 m/s in FIGURE 8b. In both speed categories, SALI improves the Doppler and PR measurements, while the gain against conventional tracking loops is larger with high speeds. This analysis is repeated for high and low C/N0 values. FIGURE 8c and 8d show a generally more improved Doppler and PR estimation for C/N0 in (20, 30)dBHz against (30, 40)dBHz. This can be explained by generally

severer MP interference associated with lower C/N0, where SALI’s MP mitigation stands out more effectively, especially for the Doppler estimation.

In all measurement epochs throughout the entire test campaign, a SALI RX either delivers a measurement with a binary MP warning flag or suppresses the instantaneous measurement due to a raised NLOS flag. For the indicated NLOS epochs, the CDFs of PR and Doppler errors associated with the conventional SF RX (ref.) are presented in FIGURE 9a and 9b, respectively.

The PR errors of the reference RX are almost uniformly distributed, and the Doppler errors’ variance is significant, showing the reliability of the NLOS exclusion by SALI. In the remaining epochs, the NLOS flag remains false, and measurements are augmented by the MP warning flag. FIGURE 9c shows the CDFs of PR errors in a SALI receiver, with and without the MP warning flag raised. The sharpness contrast of both curves indicates an effective MP warning. Moreover, the warning flag based on the MP indicator is not restricted to binary values; it can be real-valued. This would allow for soft-decisions in PE. All these features benefit the satellite selection and weighting mechanism in the navigation filter, and enhance the positioning performance indirectly.




Figure 9: NLOS flag and MP indicator performance across all test drives

POSITION ERROR, INTEGRITY AND AVAILABILITY STATISTICS IN ROAD TESTS

The measurement and horizontal position errors are all calculated against the truth logged during the test drive. FIGURE 10 shows significantly reduced (sharper CDFs) PR, LOS velocity (Doppler) and horizontal errors by SALI in all MP environments, compared to those in a conventional SF RX. The mountain plot of the CDFs reveals a substantially improved tail probability and

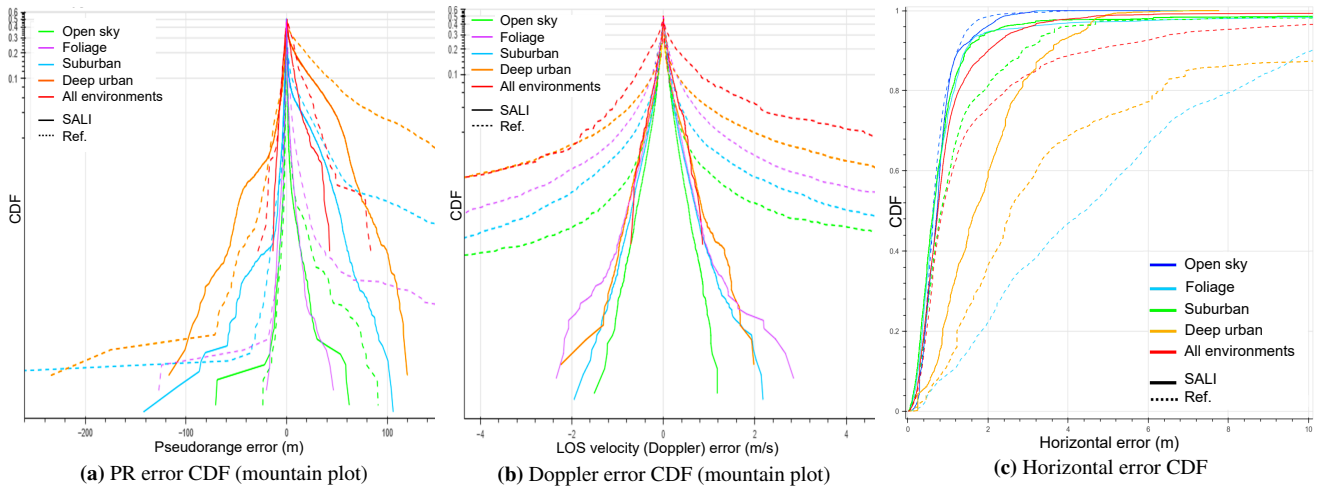


Figure 10: Measurement and position errors in different environments across all test drives


an improved zero-mean property of the error distribution, especially with the Doppler. This indicates the effectiveness of the MP mitigation by SALI. In addition, a series of safety critical metrics are defined based on the absolute value, ε , of the horizontal position error, its estimate, $\hat{\varepsilon}$, and the accuracy requirement ψ , which is two meters in our study. The statistics are aggregated over all epochs throughout all test drives, classified for different environments, and presented in TABLE 1. Here, P95 and P68 stand for the 95% and 68% percentile horizontal position error in meters, respectively. The former satisfies $\mathbb{P}(\varepsilon \leq P95) = 95\%$, and the latter satisfies $\mathbb{P}(\varepsilon \leq P68) = 68\%$. $P2m = \mathbb{P}(\varepsilon \leq \psi)$ represents the percentile of horizontal errors $\leq 2m$. The unavailability and availability probability are defined by $P_U = \mathbb{P}(\hat{\varepsilon} > \psi)$ and $P_A = 1 - P_U$, respectively. The Hazardous Misleading Information (HMI) probability is defined by $HMI = \mathbb{P}(\varepsilon > \psi | \hat{\varepsilon} \leq \psi)$. Its complement, $1 - HMI$, represents the operational probability. All statistics in MP environments are improved by SALI. The most significant gain shows up in the foliage environment as explained previously, while position error statistics in open sky conditions are marginally degraded due to SALI’s sensitivity to high order dynamics. The overall operational probability is improved from 85.2% to 94.3%, as shown by the integrity diagrams in FIGURE 11a.

Table 1: Position error, unavailability and integrity statistics across all test drives

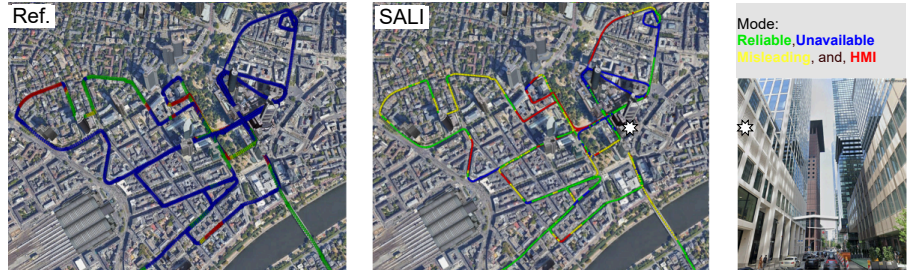
	P68 (m)		P95 (m)		P2m (%)		Unavailability (%)		HMI (%)	
	Ref.	SALI(±)	Ref.	SALI(±)	Ref.	SALI(±)	Ref.	SALI(±)	Ref.	SALI(±)
Open sky	0.8	0.9(+0.1)	1.5	1.8(+0.3)	99	96(-3)	0.0	0.3(+0.3)	1.5	3.8(+2.3)
Foliage	6.1	0.9(-5.2)	14.5	2.1(-12.4)	22	95(+73)	6.8	0.2(-6.6)	76.2	5.3(-70.9)
Suburban	1.2	0.9(-0.3)	3.7	2.1(-1.6)	82	95(+13)	20.3	7.7(-12.6)	8.3	2.6(-5.7)
Deep urban	3.3	2.1(-1.2)	15.0	4.2(-10.8)	42	63(+21)	52.1	44.5(-7.6)	34.7	25.3(-9.4)
All environments	1.4	1.0(-0.4)	7.1	2.7(-4.4)	76	90(+14)	16.9	10.7(-6.2)	14.9	5.7(-9.2)

FIGURE 11b shows a road test example in downtown Frankfurt where the operational probability is significantly enhanced by SALI, because both ε and $\hat{\varepsilon}$ fall below 2 meters in most epochs. Here, we further classify the operational epochs satisfying $\varepsilon \leq \psi$ and $\hat{\varepsilon} \leq \psi$ into two distinct modes: either the reliable mode (green) when $\varepsilon \leq \hat{\varepsilon}$ is met, or the misleading mode (yellow) when $\varepsilon > \hat{\varepsilon}$ applies.

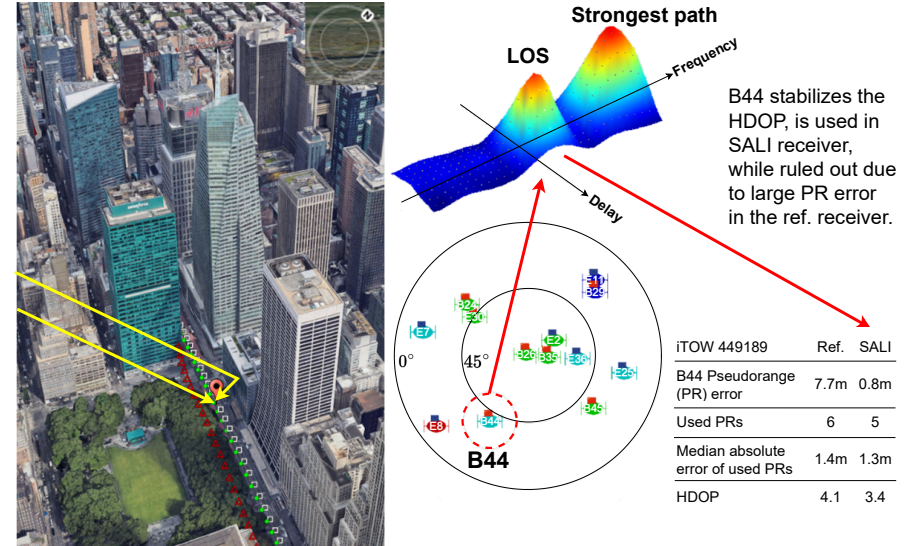
FIGURE 11c showcases the benefit of SALI with a road test example in NYC, from the PE perspective. Here, upon exiting the urban canyon, the visibility of the constellation, which includes Beidou B1C and Galileo E1, has not fully recovered. On one side of the street, there are still high-rise clusters, and on the other side, foliage introduces a significant penetration loss (and MPs). Consequently, the Horizontal Dilution Of Precision (HDOP) lacks sufficient lateral constraint. In such a scenario, satellites that contribute to the lateral constraint become highly valuable to the positioning solution. Beidou B44 with an elevation angle of about 45° would provide such a constraint if its LOS measurement is regarded as trustworthy by the navigation filter. The challenge is that apparently (at least) two impinging waves reach the antenna with minimal difference in delay and Doppler. Even more challenging is the observation that the LOS path is not the strongest one because foliage penetration loss seemingly exceeds the reflection loss by the building. The conventional RX cannot achieve such a measurement accuracy for B44 to benefit from it in this area, while SALI can resolve the LOS from the MPs, and reduce the PR error from 7.7m to 0.8m. As a result, B44 is included into the position solution. It mitigates the accumulated sensor errors, even when the SALI RX uses one PR measurement less than the reference RX. Consequently, the position fix (green dots) approaches the truth (white squares) more closely compared to the conventional SF fix (red triangles).



(a) Integrity diagrams across all test drives, SALI v.s. ref.



(b) Road test example: Downtown Frankfurt am Main, SALI v.s. ref.



(c) Road test example: NYC, B44 B1C tracking, SALI v.s. ref.

Figure 11: Road test results

The conventional RX cannot achieve such a measurement accuracy for B44 to benefit from it in this area, while SALI can resolve the LOS from the MPs, and reduce the PR error from 7.7m to 0.8m. As a result, B44 is included into the position solution. It mitigates the accumulated sensor errors, even when the SALI RX uses one PR measurement less than the reference RX. Consequently, the position fix (green dots) approaches the truth (white squares) more closely compared to the conventional SF fix (red triangles).

COMPUTATIONAL AND MEMORY COMPLEXITIES

FIGURE 12 shows the real-time CPU load measurement of a SALI-based SF RX against a conventional SF RX running on an identical hardware platform. The measurement spans from RX power-on to steady-state tracking mode, during which both RXs track an identical number of satellites. Here, the red ribbon specifically indicates the CPU load of MEs running on an ARM Cortex M3. The transition from the non-SALI to SALI mode, triggered by the successful completion of the initial sensor calibration, marks the tracking mode switch from $\leq 20\text{ms}$ CI to $\sim 1\text{s}$ SALI. Clearly, the total computational complexity of the SALI-based PE and ME decreases significantly. The main reason for the load reduction is the discontinuation of high-rate carrier/code tracking loops in each tracking channel, thanks to IMU-aided signal tracking. Unlike conventional tracking loops, SALI tends to have a larger memory footprint due to the need to store complex samples until they are evaluated by the estimator. This scales up linearly with the CI length and the number of tracking channels, in general.




Figure 12: CPU load comparison on ARM Cortex M3

of satellites. Here, the red ribbon specifically indicates the CPU load of MEs running on an ARM Cortex M3. The transition from the non-SALI to SALI mode, triggered by the successful completion of the initial sensor calibration, marks the tracking mode switch from $\leq 20\text{ms}$ CI to $\sim 1\text{s}$ SALI. Clearly, the total computational complexity of the SALI-based PE and ME decreases significantly. The main reason for the load reduction is the discontinuation of high-rate carrier/code tracking loops in each tracking channel, thanks to IMU-aided signal tracking. Unlike conventional tracking loops, SALI tends to have a larger memory footprint due to the need to store complex samples until they are evaluated by the estimator. This scales up linearly with the CI length and the number of tracking channels, in general.

DEGENERATED SALI IN ROAD TESTS

Here, we present the performance of 100ms CI coupled with 1Hz navigation rate in two cases; 1) when IMU fails, an SF RX degrades to a standalone RX; 2) SF in the navigation filter remains intact while the IMU-aided motion compensation is unavailable in ME.

For the first case, a standalone RX is tested in Hakone (Japan), a hilly area with dense foliage coverage. The test vehicle records live signals in a roundtrip, starting from the northernmost point on the map, proceeding to the southernmost point, and then returning along the same route to the northern point. On one side of the road, a continuous concrete reinforcement wall lines the path. The driving speed is brisk in general, and traffic remains minimal. Both horizontal and vertical positions are plotted in the upper and lower part of FIGURE 13, respectively. Results of 100ms CI are compared to the results of reference RXs running conventional tracking loops based on $\leq 20\text{ms}$ CI. Both are tested with eleven identical RXs connected to the same replay box. This setup allows us to demonstrate the consistency of the algorithm. A significant improvement in both horizontal and vertical position accuracy is observed with 100ms CI.

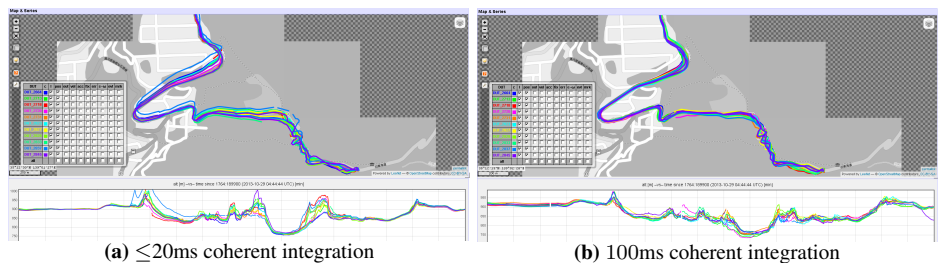


Figure 13: Road tests in Hakone, standalone RX without IMU

For the second case, an SF RX is tested in London. Here, the SF navigation filter evaluates measurements from either conventional tracking loops running $\leq 20\text{ms}$ CI (FIGURE 14a) or from the degenerated SALI with 100ms CI (FIGURE 14b). No velocity feedback is provided for motion compensation. Either scheme is tested on six RXs to show statistical fluctuation. Certainly, most of the MP damage can already be absorbed by SF. However, once again, it becomes apparent that SF algorithms gain advantages from the 100ms CI. The enhancement in position accuracy is especially evident on street curves and within urban canyon areas, where there is a marked reduction in outlier occurrences, such as the detour by DUT-4482 depicted in the center of FIGURE 14a.

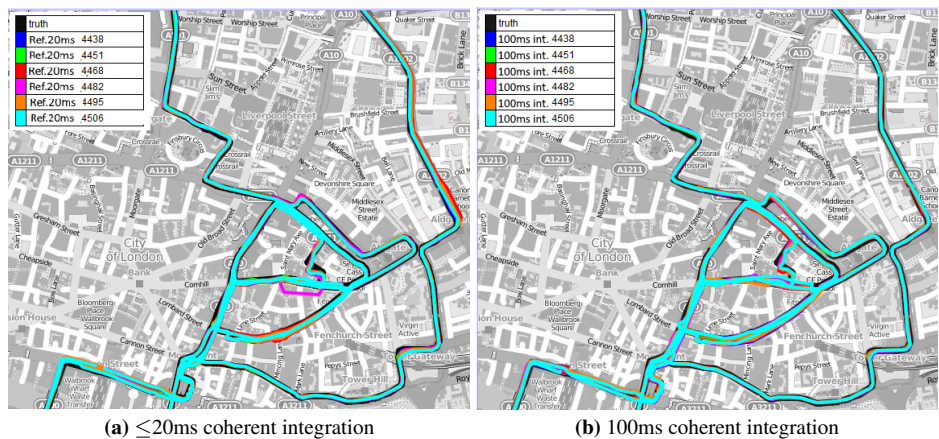


Figure 14: Road tests in London, SF RX without sensor-aided motion compensation in ME

CONCLUSIONS

We have derived a SALI-based LOS estimation scheme from a multipath signal model including both LOS geometry and stochastic channel characteristics. The proposed scheme simplifies the SAGE algorithm to just one iteration, exclusively in the LOS parameter subspace, using the IMU-aided LOS estimates as initial values. This greatly reduces the complexity and enhances the robustness. The achieved multipath mitigation results from two aspects. First, long coherent integration of ~ 1 s, realized by the IMU-aided signal tracking followed by the Gabor transform, provides fine spectral resolution, enabling separation and suppression of the multipaths, leveraging differences of multipath signatures to the LOS. As a result, multipath interference is either confined within specific frequency bins or suppressed by integrating beyond their coherence time due to multipath fading effects or Doppler rates. Second, any multipath interference falling onto indiscernible frequencies of the LOS is mitigated by the proposed delay interpolator in one step. For channels with minimal Doppler spread, the interpolator approaches the performance of the double-delta DLL without introducing correlation to the previous measurements, see FIGURE 5. The de-correlating property is especially advantageous for integrity calculations. We demonstrated enhanced measurement quality in all challenging environments in road tests, where a higher performance gain is observed with higher rover speeds and lower C/N0 values, see FIGURE 8. The most remarkable performance improvements show up in Rician fading with small K-factors, of which driving on a tree-lined (foliage) street is an example, see TABLE 1 and FIGURE 10. Moreover, the proposed NLOS exclusion methods and the multipath indicator exhibit high reliability in road tests, see FIGURE 9. The PE/ME close-loop is robust against error propagation in real-world conditions. All the performance improvements have been achieved with a substantially lower total computational complexity, see FIGURE 12, compared to a conventional sensor fusion RX running on the identical hardware platform. The backup solution with the degenerated SALI using 100ms coherent integration without IMU-aided motion compensation also exhibits considerable multipath mitigating capability, see FIGURE 13 and 14. In following situations, SALI is not expected to outperform conventional tracking loops w.r.t. position accuracy: 1) specular reflections with hardly discernible delay and Doppler to LOS, e.g. ground reflections, 2) quasi-stationary receiving conditions with minimal Doppler spread, and 3) open-sky conditions, free of multipaths, yet characterized by considerable high order dynamics. The latter can reveal the otherwise neglected error due to unknown Doppler error rates and delay interpolation, see TABLE 1.

Above all, we have demonstrated that the enhanced measurement quality by SALI translates into significant improvements in position accuracy, integrity and availability, on top of the legacy sensor fusion architecture. This ensures an enhancement of the sensor-aided high-rate positioning solution in e.g. 100Hz. We have showcased the effectiveness and robustness of SALI in a real-time GNSS RX operating in realistic conditions of frequency and delay dispersion. Such conditions are common in systems like vehicle navigation and wearable device positioning.

OUTLOOK

In this paper, all results and performance improvements stem from using a subset of L1 constellations exclusively. We anticipate additional significant benefits when processing L5 wideband signals for the same underlying physics. The frequency domain's multipath separation and suppression will enhance delay estimation at a finer scale for L5, thanks to the upgraded time resolution by the tenfold signal bandwidth. This implies an accuracy boost by SALI on top of L5's inherent multipath resistance. Consequently, we expect positioning performance to meet more stringent requirements than two-meter accuracy, such as decimeter-level precision, even without carrier phase measurements in challenging environments. The concept of a small sensor-aided hypothesis window also suggests inherent advantages in applications related to anti-spoofing and anti-jamming, akin to the benefits observed in excluding NLOS measurements.

ACKNOWLEDGEMENTS

We thank Dr. Said Moridi and Dr. Olivier Julien, as well as anonymous reviewers, for their valuable feedback.

REFERENCES

- Betz, J. W. and Kolodziejcki, K. R. (2009). Generalized Theory of Code Tracking with an Early-Late Discriminator Part II: Noncoherent Processing and Numerical Results. *IEEE Transactions on Aerospace and Electronic Systems*, 45(4):1557–1564.
- Broumandan, A., Nielsen, J., and Lachapelle, G. (2011). Coherent Integration Time Limits for Mobile GNSS Receivers in Multipath Fading Environments. *Inside GNSS Magazine*, pages 62–65. <https://insidegnss.com/wp-content/uploads/2018/01/marapr11-Broumandan.pdf>.
- Cheng, Z. (2016). Global Navigation Satellite System (GNSS) Signal Tracking. Patent: US11513235, filed Jul. 22, 2016.
- Cheng, Z. (2017). Global Navigation Satellite System (GNSS) Signal Tracking. Patent: US11187810, filed Mar. 31, 2017.
- Cheng, Z., Giger, K., and Bryant, R. C. (2018). Global Navigation Satellite System (GNSS) Multipath Mitigation. Patent:

US11579309, filed May 18, 2018.

- Faragher, R., Couronneau, N., Powe, M., Esteves, P., Crockett, M., Martin, H., Ziglioli, E., Higgins, C., and Buckle, D. (2018). Supercorrelation: Enhancing the Accuracy and Sensitivity of Consumer GNSS Receivers with a DSP Upgrade. In *Proc. of the 31st International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2018)*, pages 357–375, Miami, FL. <https://api.semanticscholar.org/CorpusID:69988198>.
- Faragher, R., Powe, M., Esteves, P., Couronneau, N., Crockett, M., Martin, H., Ziglioli, E., and Higgins, C. (2019). Supercorrelation as a Service: S-GNSS Upgrades for Smartdevices. In *Proc. of the 32nd International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2019)*, pages 256–274, Miami, FL.
- Feder, M. and Weinstein, E. (1988). Parameter Estimation of Superimposed Signals using the EM Algorithm. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 36(4):477–489.
- Fleury, B., Tschudin, M., Heddergott, R., Dahlhaus, D., and Ingeman Pedersen, K. (1999). Channel Parameter Estimation in Mobile Radio Environments using the SAGE Algorithm. *IEEE Journal on Selected Areas in Communications*, 17(3):434–450.
- Gao, G. and Lachapelle, G. (2008). A Novel Architecture for Ultra-Tight HSGPS-INS integration. *Journal of Global Positioning Systems*, 7:46–61.
- Garin, L. J., van Diggelen, F., and Rousseau, J.-M. (1996). Strobe & Edge Correlator Multipath Mitigation for Code. In *Proc. of the 9th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GPS 1996)*, Kansas City, MO, pages 657–664. <https://api.semanticscholar.org/CorpusID:59696276>.
- Groves, P., Zhong, Q., Faragher, R., and Esteves, P. (2020). Combining Inertially-aided Extended Coherent Integration (Supercorrelation) with 3D-Mapping-Aided GNSS. In *Proc. of the 32nd International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2020)*, pages 2327–2346, Virtual conference.
- ITU-R P.681-8 (2015). Propagation data required for the design of Earth-space land mobile telecommunication systems. Standard, International Telecommunication Union (ITU), Geneva, CH.
- Lehner, A. (2007). *Multipath Channel Modelling for Satellite Navigation Systems - Mehrwegekanalmodellierung für Satelliten-navigationsysteme*. PhD thesis, University Erlangen and DLR.
- Lentmaier, M., Krach, B., and Robertson, P. (2008). Bayesian Time Delay Estimation of GNSS Signals in Dynamic Multipath Environments. *International Journal of Navigation and Observation*, 2008.
- Luo, Y., Hsu, L.-T., and El-Sheimy, N. (2023). A Baseband MLE for Snapshot GNSS Receiver Using Super-Long-Coherent Correlation in a Fractional Fourier Domain. *NAVIGATION: Journal of the Institute of Navigation*, 70(3). <https://navi.ion.org/content/70/3/navi.588>.
- McGraw, G. A. and Braasch, M. S. (1999). GNSS multipath mitigation using gated and high resolution correlator concepts. In *Proc. of the 1999 National Technical Meeting of The Institute of Navigation, San Diego, CA*, pages 333–342.
- Pany, T., Riedl, B., and Winkel, J. (2009). Coherent Integration Time: The longer, the Better. *Inside GNSS Magazine*. <https://insidegnss.com/auto/novdec09-wp.pdf>.
- Schmidt, G. T. and Phillips, R. E. (2010). INS/GPS Integration Architectures. <https://api.semanticscholar.org/CorpusID:18288801/>.
- Schmidt, G. T. and Phillips, R. E. (2011). INS/GPS Integration Architecture Performance Comparisons. <https://api.semanticscholar.org/CorpusID:18821037/>.
- Somieski, A., Hollenstein, C., Favey, E., and Schmid, C. (2010). Low-Cost Sensor Fusion Dead Reckoning using a Single-Frequency GNSS Receiver Combined with Gyroscope and Wheel Tick Measurements. In *Proc. of the 23rd International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS 2010)*, pages 1645–1652, Portland, OR. <https://api.semanticscholar.org/CorpusID:60539481>.
- Van Dierendonck, A. J., Fenton, P., and Ford, T. (1992). Theory and performance of narrow correlator spacing in a gps receiver. *NAVIGATION: Journal of the Institute of Navigation, Volume 39, Number 3*, 39(3):265–284.
- Van Nee, R. (1992). The Multipath Estimating Delay Lock Loop. In *IEEE Second International Symposium on Spread Spectrum Techniques and Applications*, pages 39–42.
- Xie, P. (2023). Multipath Mitigation. *Inside GNSS Magazine*. <https://insidegnss.com/multipath-mitigation/>.